THORN EMI Software Software Sciences

AIDA

FINAL REPORT

19980309 252

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Accession Number: 3539

Publication Date: Aug 01, 1990

Title: AIDA Model 1.0 Final Report

Corporate Author Or Publisher: Software Sciences Limited, Famborough, Hampshire, England

Comments on Document: Final Report

Descriptors, Keywords: AIDA UK Hardware Software UKAS Operational Requirement Theater Missile Defense TMD Weapon Sensor Battle Management Command Control SDS

Pages: 00023

Cataloged Date: Jun 03, 1992

Document Type: HC

Number of Copies In Library: 000001

Record ID: 24003

AIDA

FINAL REPORT

Date
August1990

Software Sciences Limited Farnborough Hampshire England GU14 7NB
telephone (0252) 544321 telex 858228 fax (0252) 513739

Project:	AIDA	MOD Contract: SLS21C/1	
Project:	Development of a Prototype Artificial Intelligence Framework for Theatre Missile Defence		
Manual:	Design	Volume: 26	Part: 1
Author:	D.R.Sloggett C.J.Slim	Date: 27.7.90	Issue: 2
Title:	AIDA Model 1.0 Final Report		
Doc. Ref.:	AIDA/3/26/01		
Cross Ref.:	AIDA/1/06/01 Quality Plan		

Contents:

1. Introduction

1.1 Overview

This report contains a summary of the main activities and results of the first year of the Artificial Intelligence Discrimination Architecture (AIDA) programme. The programme arose from a joint US/UK initiative - outlined in a white paper from the Special Systems Department (SS) in the Royal Aerospace Establishment (RAE) Farnborough (Reference 1) - to follow on research into the subject of discrimination initiated under the United Kingdom Architecture Study (UKAS) programme.

The main aim of the project was to develop a Framework into which discrimination algorithms, simulated and real threat data, and sensor models representing real world elements of a Ballistic Missile Defence (BMD) system could be placed. The Framework would provide facilities to allow threat locations, raid composition, flight dynamics, sensor viewing and sensor detection processes to be modelled. Once signature data was extracted from the sensors facilities to process and analyse the discrimination data and present the results of this analysis to an operator were also required.

In parallel with the development of the Framework the programme provided for continued research into the topic of discrimination, notably in the mid course phase, using innovative algorithmic concepts. The aim being to define a comprehensive TMD discrimination algorithm. In the first year of the project this research has studied signature analysis algorithms in the boost and mid course phases and has defined an infrastructure in which these algorithms can operate as part of a comprehensive end-to-end classification and discrimination process within a TMD system.

This report reviews the background to the AIDA programme and highlights key results arising from the first years research activities. The report also reviews the structure and progress made on the development of the experimental Framework.

The AIDA programme was started on the 23rd of January 1989. In its first year the main goals of the programme were to:

- Develop and document a Framework in which the application of AI techniques to discrimination within TMD and BMD architectures could be evaluated.

- Continue research into innovative discrimination concepts and the definition of the structure and operation of an operational discrimination algorithm for a TMD system.

- Report on the architectural issues arising from such an approach to discrimination.

- Compile an evaluation programme designed to highlight the utility of the Framework and algorithms in the context of a European TMD system.

- Compile a User Guide that would allow third parties to use the system for their own evaluations.

- Compile a plan recommending additional research on the framework.

1.2 The UKAS Programme and AIDA

The UKAS programme is a joint US/UK research programme conducting research to establish outline Operational Requirements (OR) for a Theatre Missile Defence (TMD) system. The aim of the UKAS programme is to define an architecture for a TMD system that comprises weapons, sensors and battle management, command and control functions. The programme is currently in its seventh phase of research where studies have been conducted into the potential defence of Europe that a US Phase 1 Strategic Defence System (SDS-1) could provide.

In the course of the UKAS research the requirements for the performance of elements of the architecture have emerged from detailed studies of the potential ballistic missile threat to Europe, definition of the mission for a TMD system and the assets such a system should defend - with its coupling to deterrence - and technology studies aimed at defining those technologies that are feasible within a number of defined epochs for deployment of a system (Reference 2). These epochs concentrate on the mid term (2005-2015) with studies of the near term (1995-2005) and far term (2015-2025) being carried out as excursions on any mid term architecture. UKAS research was initially aimed at the era preceding the Intermediate Nuclear Forces (INF) Treaty. Phases 1 to 4 of the programme studied the technologies of a mid term architecture aimed at defending Europe from a massive pre-emptive attack by Soviet Intermediate Range Nuclear Forces (IRBM) and Short Range Nuclear Forces (SRBM) such as the SS-20 and the SS-23. Subsequent phases of the programme have concentrated upon the post-INF Treaty era where revised threat, mission and technology studies have updated the UKAS programme scenarios and architecture requirements.

1.3 The Basis of the AIDA Programme

In the course of the UKAS programme the topic of discrimination has been identified as one of the critical cost drivers for the system. The results of the Phase 6 architecture studies highlighted this issue (Reference 3).

Studies carried out by SS RAE and Hunting Engineering Limited (HEL) defined a threat for the UKAS that represented UK state of the art thinking at that time (Reference 4). This threat was analysed by the UKAS sensor technology team to determine if any discrimination data of value to the defence could be derived from the threat. This analysis showed that the threat designers had removed all of the potential discriminants from the threat. The architecture design team was therefore forced - in the absence of any advanced sensor technologies, such as interactive sensors - to devise an architecture that fired at each threat object detected and tracked by the defence. Whilst it is recognised that this provided a somewhat exaggerated solution to the threat it nevertheless provided an important insight to the upper bounds of feasibility in respect of TMD architecture analysis.

The threat proposed by SS RAE and HEL comprised some 80,000 objects in a raid upon Europe. The cost of the mid course weapon system whose task was to engage all of these false targets accounted for over 65% of the total cost of the architecture. The total estimated cost of the European architecture being greater than the total cost of the US SDS-1 system. This graphically illustrated the linkage between discrimination and the overall cost of a TMD architecture.

In recognising the importance of discrimination as a function within a TMD system a proposal (Reference 5) was made by Software Sciences Limited (SSL) through SS RAE to the US Strategic Defence Initiative Office (SDIO) in Washington to explore the topic in far greater detail than had been possible in the course of the UKAS programme. The proposal envisaged the development of a Framework based upon Artificial Intelligence (AI) techniques in which high fidelity signatures and algorithms could be evaluated for all flight regimes in a TMD system. The aim was that the Framework should be able to support modelling of both TMD and BMD sensors and allow detailed analysis of threat signatures to be carried out as a simulation of the real time operation of a TMD or BMD system. The proposal compiled by SSL was consolidated by SS RAE and SDIO into Reference 1 to act as a formal statement of work for the AIDA Phase 1 programme.

1.4 The AIDA Research Goals and Objectives

The emphasis in the AIDA programme was on the continuing research and development of new and innovative approaches to discrimination and their examination and proof of utility in the presence of a variety of countermeasures on the Framework.

In the course of the UKAS programme research had highlighted the possibility of using feature analysis techniques to extract key features from signature space that would allow the defence to optimise its engagement philosophy in the mid course phase. The concept was based upon real time analysis of signature data derived by the sensor systems in the architecture and not to place any significant reliance upon a-priori threat data gained through intelligence channels (Reference 6).

The adaptive algorithms would attempt to identify and learn the location in feature space of the warhead. Analysis of objects where countermeasures had failed to deploy correctly was seen to be one of the key elements of this approach.

In processing the features derived by the discrimination algorithms attempts would be made to identify small groups or clusters of objects that exhibited similar characteristics to warheads. Where objects were seen to group together hypotheses would be established in an attempt to match the observed characteristics to descriptions of warheads. Those smaller groups achieving a match providing important clues as to the underlying structure within feature space and the location of the warheads.

This identification of the location in feature space of the warheads - in direct contrast to their positive identification through the extraction and analysis of some clearly defined, and in some cases already known discriminant - provides a different and much more realistic approach to the analysis of discrimination data. The concept is known as the analysis of **partial discriminants** or the **Failure analysis algorithm** and is at the core of the AIDA research programme (Reference 7).

In the course of the phase 1 programme the concept of analysis of partial discriminants has been further defined. A classified annex accompanying this report outlines the main theoretical research findings from the first phase of this work and the definition of a comprehensive discrimination algorithm for an operational TMD system. The features of this algorithm include:

- A real time AI based approach to Boost phase missile classification
- A real time AI based approach to Mid Course discrimination
- Real time adaptation of feature space analysis to the threat detected by sensor systems
- Limited reliance upon a-priori data
- A hierarchy of clustering algorithms designed to adaptively analyse the structure of feature space
- The practical application of partial discrimination analysis to an operational TMD system
- An ability to learn the structure of the threat in real time.
- Analysis of feature space data in the context of the defence assessed raid and mission objectives.

2 The AIDA Phase 1 Programme

2.1 Overview

This section highlights the key events in the AIDA Phase 1 programme and describes the equipment used on the project and the effort placed on the project by SS RAE.

2.2 Project Hardware and Software

2.2.1 Project Hardware Configuration

The programme was to be based upon a single Sun 3/60 workstation based at SSL. Early on in the programme it was realised that this machine would limit the development programme for the software and SS RAE agreed to provide additional machines purchased from intramural funds. Two additional Sun workstations, a line printer and tape drive were obtained by SS RAE to support the project and expand their computing facilities. This additional level of support was critical to the successful outcome of the project and represents a significant additional contribution to the programme by SS RAE.

2.2.2 Project Software Configuration

The AI element of the programme was to be based upon the Intellicorp Knowledge Engineering Environment (KEE) product. Studies carried out at the start of the programme showed this to be a more flexible tool than the NEXPERT tool originally proposed at the start of the project.

The US required that as much of the software as possible be developed in ADA. This necessitated the procurement of an ADA compiler for the Sun workstations. After a survey of the potential suppliers SSL concluded that the product offered by ALSYS would be the best supported and decided to purchase their software package. Unfortunately it soon became clear that this product contained some deficiencies that prevented use of ADA in those areas external to KEE where the intention was to code the software in ADA. Solutions based upon the language 'C' and a graphics package called 'Wave' were found to offer a cost effective way forward. SS RAE purchased the 'Wave' package from MOD Pats 3C on behalf of the project.

2.3 The Stages of the Phase 1 Programme

2.3.1 The Design Phase

The design phase was an activity scheduled to last for 4 months at the start of the programme. In the course of this phase the scope of the design work embraced all flight regimes of ballistic missiles. At a Critical Design Review (CDR) meeting held at SSL at the end of this phase it was agreed that the design work should focus on the boost and mid course phases of flight and that the post boost and terminal phases would be left until a future date. In this phase two In Progress Review (IPR) meetings were scheduled. In the course of the first review meeting SSL presented their outline ideas for modelling all phases of the flight regimes, and their underlying requirements analysis results defining assumptions for the Framework. The second meeting comprised the CDR.

The main design review meeting was then rescheduled for 2 months later at which point agreement to proceed to coding and developing the Model 1 system - based upon boost and mid course simulations - was received.

In the course of this phase of the project a survey of clustering algorithms was carried out in order to find a method of clustering that would integrate within the discrimination concept. The results of this study were reported in Reference 8.

2.3.2 The Coding Phase

The coding phase started at the end of June 1989 and lasted until the beginning of December 1989. No distinct phase of module integration was possible as the complexity of the code varied between modules and some were produced in September and others were delayed until the end of December. Demonstrations of the Framework were made at IPRs in September and December to show progress on the project.

2.3.3 The Module Testing and Integration Phase

This phase of the project lasted from the end of December until the start of the first scheduled acceptance tests at the end of January 1990. An IPR was scheduled to review the outcome of this acceptance test and was held at the end of January. At this time the Framework as envisaged at the start of the coding phase had not been completed and SS RAE agreed to support a 3 month extension to complete the full scope of the programme. This additional 3 month effort has been supported by RAE and resulted in a final acceptance test being carried out at the end of April 1990 by both SS RAE and SDIO. The Model 1 system has subsequently been delivered to SDIO in Washington.

2.3.4 The System Acceptance Phase

The system acceptance phase culminated in a three day acceptance test programme being run at SS RAE and SSL at the end of April 1990. In these tests, Model 1.0 was validated against the agreed Acceptance Test Plan (Reference 12). The tests comprised a series of short runs to prove framework integration and detailed module tests to demonstrate their functionality and integrity. These tests also provided initial performance measures for the AIDA Model 1.0 system. The results from the baseline runs performed as part of the acceptance tests are summarised in Section 4 of this paper.

On the basis of these tests the model was accepted as functionally complete and fully integrated. A copy of the model was duly despatched to SDIO.

2.4 SS RAE Support to the Programme

2.4.1 Data

In the course of the phase 1 programme RAE supplied a number of items of data to the project team to initialise the threat database for boost phase and mid course signatures.

2.4.2 Labour

In the course of the project SS RAE provided invaluable design expertise to the project and monitored the programme on a monthly basis. In addition to the management support the project received advice and assistance from many SS RAE experts in the field of discrimination.

In addition to the expert advice SS RAE directly supported the programme with the involvement of one member of RAE staff whose task was to design, code and test the raid generation facilities.

2.4.3 Computing Resources

The two additional SUN workstations and associated software (including the Wave and KEE packages) provided by SS RAE (see 2.2) assured the successful completion of the project.

3. The AIDA Framework

3.1 Overview

This section summarises the design and current development statue of the Model 1 AIDA Framework. The outline concepts and design ideas have already been documented in Reference 9.

The AIDA Framework comprises 3 main subsystems: the Scenario Generation (SG) subsystem; the AI Discrimination subsystem (AID) and the Performance Assessment (PA) subsystem. The SG subsystem is designed to provide facilities whereby a operator - through the Man Machine Interface (MMI) - can initialise threat and defence sensor parameters. A full description of the operation of this facility can be found in Reference 11.

In the course of defining the requirements for this framework a number of assumptions were made to limit the scope of the AIDA Model 1 system. These are documented in Reference 10 but are re-stated here for completeness.

The AIDA Model 1 system is based upon the following design assumptions:

- Only fixed sensors would be modelled.
- No nuclear effects would be modelled
- No tracking would be modelled
- The loading on sensor systems would not be controlled
- No star backgrounds for mid course

3.2 The Scenario Generation Subsystem

The SG subsystem of the AIDA Model 1 Framework can be divided into the following parts: the MMI module, the raid generation module, the sensor viewing module, and signature generation and validation modules for the boost and mid course phases of flight.

The MMI allows for an operator to specify the component parts and assets under attack in a raid. This raid can then be generated through a process of modelling in high fidelity the boost, post boost and mid course phases of flight of the missiles, warheads and decoys. This raid generation process generates a file containing a summary of the state vectors of each item in the raid as a function of time. Once created the raid file is read by the sensor viewing module. This module computes the time in the raid at which it is possible for the sensor systems defined in the defence architecture to view threat objects. The basis of this algorithm is that objects should be within the Field Of Regard (FOR) of the sensor system, i.e. over any earth horizon , but not necessarily detected by the sensor system - this function being calculated in real time as a function of aspect angle, body motion and range effects.

During the time when it is run the sensor viewing module generates a file containing the time at which objects enter the FOR of each sensor in the architecture. This file is used by the signature generation modules to determine when a signature should be generated for each threat object across the interface with the AID subsystem.

The task of the signature validation modules is to provide a facility whereby operators can display the signatures generated in a specific scenario. This allows signatures to be validated before the AID subsystem is activated.

3.3 The AI Discrimination Subsystem

The AID subsystem comprises three levels of processing carried out on the signature and threat data. In the level 1 processing the signature data generated from the SG subsystem is received - either in run time on a time stepped basis or from a pre-stored file - by modules simulating defence sensor systems. The level 1 modules - known as Knowledge Level 1 (KL-1) - simulate the attribute extraction process that would be undertaken by defence sensor systems at source. The KL-1 modules simulating radar and IR sensor receive, statistically process and analyse attributes of each threat object being viewed by sensor systems.

Once the attribute data has been derived by the KL-1 modules this data is published to the KL-2 module. In each flight regime one KL-2 module operates in order to analyse the attributes derived by the KL-1 module and attempt to group these attributes to analyse any underlying structure in feature space that might allow discriminants to be identified and exploited.

Control of the AID is carried out by the Framework control module. This schedules the tasking of modules within the AID subsystem controlling the transfer of data from the SG system into the level 1 processes, the activation of the level 1 process, and the activation of the level 2 processes to carry out the multi-feature space analysis of the data.

During run time the Monitor facility provides a direct read out to the operator of the current results of the test run. In the boost phase the numbers of classified missiles and their identities are summarised on a colour display. In the mid course phase the monitor can display the numbers of objects classified as RVs and decoys and the structure of the feature space data on which this analysis has drawn.

Another new aspect of the AIDA research was the introduction of the idea of placing the analysis of the discrimination data in the context of the perceived mission objectives of the offence and their likely targeting doctrine. In the AIDA Model 1 system this analysis is carried out in the KL-3 raid analysis module.

The objective of the KL-3 module is to perform the analysis of the mission objectives of the offence - using Warnings and Indicators (WIs) to place such an analysis in context - and to predict the scale of the raid on the basis of the partial observations made through cloud cover in the first 50 seconds of the raid. This latter function being seen as a vital element providing early advice to TMD commanders as to the likely intentions of the offence.

Whilst not implemented in the Model 1 system the research carried out in the first year of the project has also defined the scope and operational concepts behind two additional modules required in the overall battle management hierarchy; KL-4 and KL-5. The KL-4 module is seen to take the data from the KL-3 raid analysis and KL-2 classification and discrimination analysis and perform an optimum Weapon Target Assignment (WTA) algorithm to maximise the value of the offence missiles and warheads destroyed in a specific layer of the defence.

The role of the KL-5 module would be to accept defence mission objectives from a TMD commander and to interpret an optimum response to the threat analysed by KL-2 and KL-3. In this role KL-5 would analyse the results of intercepts in each layer and re-define the defence mission objectives before weapons were assigned by KL-4 in the next defence layer. What would flow from KL-5 to KL-4 is a revised asset protection list that would change once the outcomes from intercepts in each layer become clear. KL-4 would then be activated - once KL-3 had completed its next analysis - such as in the mid course phase, and would decide upon the optimum allocation of exo weapon systems in a first salvo of weapons; the same process being repeated for any subsequent salvos that were required.

3.4 The Performance Assessment Subsystem

The task of the PA subsystem is to display the results of specific runs on the Framework. The system MOPs and MOEs that have been defined for AIDA all relate to the numbers of missiles or warheads classified or discriminated in each phase of the simulation. The PA subsystem reads data stored on files in the course of a run and displays them to an operator once the run has been completed.

4. Phase 1 Results and Conclusions

4.1 Results

This section summarises the main results from using the Framework derived from Phase 1 of the AIDA programme. These are expanded in greater detail in the accompanying classified annex.

It is not possible as yet to make definitive statements as to the utility of the Framework in either the boost or mid course phases. Results will shortly be available from the evaluation plan where detailed experiments will be carried out to study the effects of different countermeasures on the performance of the algorithms.

Early experiments with the framework appear to show the utility of the AIDA Phase 1 algorithms in classifying Soviet ICBM and SLBM boosters in the boost phase. These evaluations have been carried out in the period leading up to the final acceptance test. They only represent insights to the effectiveness of the algorithms in the presence of countermeasures and should therefore be treated with some caution until the final evaluation phase commences. However the algorithms have been run in the presence of false staging events and have continued to identify key features in the threat data. These results are encouraging and show some resilience of the algorithms to countermeasures.

In the mid course phase simple evaluations of the utility of the Partial Discriminants approach have been hampered by low fidelity modelling of some of the threat objects. Currently the results of the clustering process appears to perform ahead of that expected. Increased fidelity in threat signature generation will show the ability of the algorithm to work in more difficult signature environments.

4.2 Lessons Learnt

The validation runs performed in the acceptance tests form a baseline for future enhancements. These tests identified that there were certain run-time constraints associated with the Phase 1 Model 1.0 system.

These run-time constraints centred on the use of the KEE (Knowledge Engineering Environment) package for implementing the discrimination algorithms. This tool provides considerable flexibility and a rapid prototype development environment which has to be traded against its run-time and memory requirements. In practical terms, the tests required 14 hours to run a 10 booster scenario through the complete boost phase and 1 hour to perform the discrimination activity for each simulation second processed in mid-course.

In addition, significant processing is required by the signature generation module (coded in the Wave package), assessed at 30 minutes to perform the generation of signatures for 500 mid-course objects from sensor readings over a single simulation second. The choice of the software package in this case is probably not as critical as the amount of processing required to perform the high fidelity modelling implemented within the signature generator.

The primary conclusion from this analysis was that one of the tasks performed within Phase 2, in advance of the evaluation programme, should involve a run-time consolidation exercise.

It was established that the MMI is sensitive to the types of data entered and that the offence database files produced are restricted in their portability. These problems are a direct consequence of the selection of Ada for the task, and specifically the Alsys compiler (see 2.2.2). Solutions require enhancements to the code, including the addition of extra facilities, or (probably the better option in the long term) the use of a commercial database package in conjunction with the existing MMI screen displays.

It has also been established that, apart from its run-time and memory limitations, KEE constitutes a useful tool for prototyping and developing AI systems. The discrimination subsystem of the AIDA system has been developed within the KEE environment. This has enabled a prototype discrimination framework to be implemented very rapidly and has permitted considerable refinement and enhancement to be undertaken - to incorporate further algorithms and to extend the intelligent processing.

KEE is an object-oriented programming environment providing the facilities for producing blackboard systems with

- distributed processing goals
- distributed domain and system knowledge
- goals which are activated as information is posted on a blackboard
- an excellent user interface for knowledge engineering and analysis purposes
- embedded LISP code - being an interpretive language it aids rapid prototyping and provides excellent debugging facilities
- a production rulesystem which incorporates both forward and backward chaining, a variety of rule search strategies and useful rule trace and debug facilities; unfortunately, rule development and testing has been found to be extremely cumbersome and slow and consequently, where practical, rules have been produced within the LISP code instead

4.3 Conclusions

Research in the UK has established that any discrimination concept based solely on the principle of a K-factor or Bayesian analysis would not provide a robust solution to warhead identification in TMD system operation. This research thus established the need for an adaptive, real time approach to discrimination in BMD architectures based upon a reduced emphasis on the use of a-priori data.

The following conclusions can be drawn from the Phase 1 AIDA programme:

- An AI Framework has been created and some initial evaluations, conducted in the process of validating and testing the facility, show promising results.

- As part of the framework, a stand-alone scenario generation subsystem has been produced to simulate the TMD environment, providing signature and track data on each object visible to the sensor systems modelled.

- Use of the AI approach has enabled rapid development of the prototype framework and has been shown to provide the flexibility for exploiting analysis based upon both deterministic and intelligent processing.

- Algorithms and reasoning have been implemented to perform classification of objects which is not dependent upon the existence or quality of a-priori knowledge, deduced from trials or intelligence sources.

- A Failure Analysis Algorithm has been devised and integrated within the Framework. Initial results do not provide conclusive proof of its utility in TMD scenarios.

- The framework produced provides facilities to model and assess the capability of a TMD architecture to detect the and analyse the components of a ballistic missile raid. No evaluation has yet been undertaken to generate any measures of architecture performance.

In addition, a Phase 2 programme has been identified (see Reference 14) whose principal goals will be:

1. To perform a consolidation task, in anticipation of an evaluation of the AIDA Model 1 system. The objectives of this exercise would be twofold:

   - identify and incorporate run-time enhancements; any reduction in run-time achieved should be benchmarked against the baseline runs

- determine the number of objects and length of scenarios that are practical to assume for the evaluation program (the target being for runs to be performed within an overnight period)

2. To enhance Model 1 to incorporate a post-boost phase model. This will be the AIDA Model 2 system and will require the design, implementation, integration and testing of functions and facilities for

   - simulating the post-boost scenario and generating the signatures for the objects associated with the post-boost phase of flight

   - development of discrimination algorithms using AI techniques to classify the objects from their signature data

   - provision of additional facilities to derive measures of effectiveness in the post-boost phase

3. Execute the evaluation program for Model 1 and Model 2 (see Reference 13), to assess the ability of the AIDA system to enhance the discrimination performance of TMD architectures.

4. Assess the overheads on a TMD system that would arise from the practical implementation of an AIDA approach to classification and discrimination.

5. Propose a plan for Phase 3 to extend Model 2 to incorporate a terminal phase and enhancements arising from the evaluation of Model 2.

6. Study and define the interface requirements necessary in order to transport the Model 2 system onto the EADTB at RSRE, Malvern.

5. AIDA Phase 1 Outputs

The main results of the first year activities on the AIDA programme can be summarised as follows:

- A Framework supporting an AI based analysis of signatures in a TMD environment has been developed.

- A Framework supporting the generation and analysis of high fidelity boost and mid course Infra-Red (IR) and radar signatures has been developed.

- The Framework elements required to support threat definition, raid generation, sensor viewing and signature validation have all been designed and coded.

- The algorithms required to support the analysis of small member classes and the extraction of partial discriminants have been designed and coded.

- A survey of clustering and classification algorithms has been completed and a report generated.

- Some preliminary runs have been made on the Framework to test and validate its operation.

- A Model 1.0 system has been delivered to SDIO in Washington.

- A User Guide describing the features and operation of the AIDA Model 1.0 system has been published (Reference 11).

- The Model 1 Framework has undergone acceptance testing by the US SDIO and its support contractors General Research Corporation (GRC) in the UK against an agreed and published acceptance test plan (Reference 12).

- A Plan for the evaluation of the Model 1 system has been published (Reference 13).

- A recommendation plan for a follow-on research programme has been published (Reference 14).

- The research into the basis of an operational TMD discrimination algorithm has continued and a report summarising the results of this work has been completed (see the classified annex to this document).

- The US and UK Theatre Missile Defence community has been briefed on the project and its results at the TMD conference in London and the US/UK Countermeasures Score Group meeting in Washington (Reference 7).

- The UKAS team has been briefed on the programme and its initial conclusions through participation in UKAS Design Nucleus (DN) meetings.

- The US AIDA management team have been briefed on the progress and plans for the project at bi-monthly IPRs held in the UK.

6. References

1. RAE SS (1988), Proposals for the Development of a AI Based Framework for Discrimination Studies.

2. E. Haysler (1988), The UKAS Phase VI Mission Requirements Analysis.

3. UKAS Team (1989) , UKAS Phase 6 Final Reports

4. R.Chase, R.Talbot (1988), The UKAS Phase V Threat Statement.

5. D.R.Sloggett (1988) , Proposals for the Development of a AI Based Framework for Discrimination Studies.

6. D.R.Sloggett, (1989), The Application of AI based techniques to Discrimination in TMD systems, MILCOMP 1989 Conference, London, England.

7. D.R.Sloggett, (1990), The Application of AI based techniques to Discrimination in TMD systems, TMD Conference, London, England.

8. C.J.Gooding (1989) , A Review of Clustering Algorithms for the AIDA Project.

9. D.R.Sloggett (1989), AIDA Framework Design Ideas

10. C.Price (1989), AIDA Requirements Analysis

11. M.Dreksler et al (1990) , AIDA Model 1.0 User Guide Ref: AIDA/4/01/01

12. C.J.Gooding (1990) , The AIDA System Acceptance Test Specification. Ref: AIDA/3/25/01.

13. C.J.Gooding (1990), The AIDA Evaluation Plan. Ref: AIDA/3/23/01.

14. C.J.Gooding (1990), The AIDA Recommendation Plan. Ref: AIDA/3/24/01